TIME-DELAYED OPERATION OF A TELEROBOT VIA GEOSYNCHRONOUS RELAY

Brian H. Wilcox
Supervisor- Robotics Systems Research,
Implementation, and Integration Group
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109
(818) 354-4625

ABSTRACT

Operation of a telerobot is compromised if a time delay more than a few hundred milliseconds exists between operator and remote manipulator. However, the most economically attractive way to perform telerobotic functions such as assembly, maintenance, and repair in Earth orbit is via geosynchronous relay satellites to a ground-based operator. This induces loop delays from one-half to two seconds, depending on how many relays are involved. Such large delays makes direct master-slave, force-reflecting teleoperated systems infeasible. Research at JPL on a useful telerobot that operates with such time delays is described.

INTRODUCTION

It has long been recognized that the performance of a master-slave teleoperator is seriously degraded if the closed-loop latency between input at the master, action at the slave, and perception of the result back at the master is more than a few hundred milliseconds [1]. It is generally felt that loop delays of a few milliseconds are acceptable, that delays of a few hundred milliseconds can be overcome with extensive training (untrained operators are unstable in their attempts to guide the slave with such delays), and that delays of 500 milliseconds or more cause operators to revert to a 'move and wait' strategy for manipulator control. The move-and-wait strategy gives very low performance compared to normal teleoperation, and is still prone to errors and damage due to undesired contact between the slave manipulator and its environment.

The space tasks of greatest interest to NASA and the Air Force are those in low Earth orbit, and of course the most economical place for the operator is on the ground. Continuous communication between these two sites is most easily accomplished via one or more geosynchronous communication satellites, at an altitude of 22,300 miles above the equator. The time delay at the speed-of-light from Earth or low orbit to geosynchronous orbit is about 120 milliseconds. Thus the most straightforward and economical arrangement of a teleoperator for space tasks has a round-trip latency of 1/2 to 1 second (actually 480 milliseconds if only one geosynchronous relay is used, and double that if line-of-sight considerations cause two relays to be needed). In fact, the delay can be as much as two seconds if the ground operator is not located at one of a few Earth-station points such as White Sands or Palo Alto, since a second satellite link would be used from the operator's location (e.g. Houston) to the prime Earth-station site. Thus the minimum latency of a geosynchronous relay will cause serious degradation of the performance of a conventional teleoperator system. (Operation by astronauts is of course possible without significant time delay, but astronaut time is very expensive, and many U.S. space assets are in orbits that cannot be reached by the Space Shuttle.)

A TIME-DELAYED TELEROBOT

Supervisory Control was proposed a decade ago to deal with this problem [2]. This paper describes a particular implementation approach for supervisory control being studied at JPL as part of the Telerobotics Research program sponsored by the NASA Office of Aeronautics and Space Technology. The architecture of this telerobotic system is as shown in figure 1.[3]

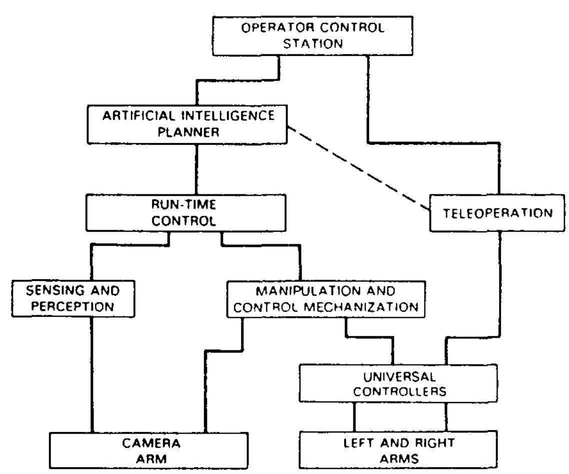

Figure 1. Telerobot Subsystem Architecture

The subsystems of this architecture which interact with the environment are the Sensing and Perception subsystem and the Manipulator Control Mechanization Subsystem (S&P and MCM) [4][5]. The actions of these subsystems are coordinated by the Run-Time Control subsystem (RTC) [6]. The RTC contains the task data base, performs fine-motion planning, requests verification of the locations and orientations of objects by the S&P subsystem, and commands trajectory via points and force/torque task frames for execution by the MCM subsystem. The Artificial Intelligence Planner (AIP) performs gross motion planning and automatic task sequencing, and gives the operator a 'window' into the more autonomous operations of the telerobot so that planned actions may be reviewed and modified [7][8]. The teleoperation channel allows force-reflecting hand controllers

(FRHC's) to directly control the manipulators, with forces and torques detected at the wrists of the manipulators and used to backdrive the FRHC's.

It is of great interest to determine which subsystems must be resident at the remote site, since of course the flight qualification of high-performance computational elements is very demanding. Ground-based computers can be extremely powerful compared to space-qualified computers, and of course weight and power consumption are practically negligible considerations on Earth. Thus an important objective of research with the telerobot testbed is to determine the minimal flight elements of a useful telerobot, given the time delay of the geosynchronous relay link.

Physically, all the subsystems of the telerobot testbed are connected via ethernet (except the teleoperation channel, which uses a high-speed parallel channel to the Universal Controllers [9]), allowing general communication protocols and easy reconfiguration. This further allows the simulation of different assignments of subsystems to the operator site or the remote site, by inclusion of a time-delay simulation in the forwarding of the ethernet packets. It has long been recognized that a simulation of system loop delay need only put the entire delay in one leg or the other of the communication system, and that it is not necessary to split the delay between uplink and downlink, as it is in a real operational system [1]. Since the downlink from the telerobot will include several channels of video, and since this will ultimately be high-quality color video (e.g. red, green, and blue transmitted separately), the bandwidth of the downlink can easily approach 1 Gigabit per second. (Of course, operational considerations may force the use of slow-scan, image compression, and other techniques to reduce this, but these expediencies are outside the scope of the current testbed focus of basic telerobotic research.) The uplink, on the other hand, will have a few tens of ethernet packets per second, at a few hundred bits each, plus data from the two FRHC's consisting of the twelve encoder values sampled at a kilohertz (our current rate, which gives very good results), for a total of a hundred kilobits per second. Thus delay simulation is most easily accomplished by buffering the uplink. We are currently designing a time-delay simulation element of the architecture which essentially 'gateways' messages on the ethernet between subsystems declared to be at the operator site and those declared to be at the remote site, as well as buffering the FRHC data. The 'downlink' ethernet packets, arm position and reflected force data, as well as video and other sensor data, are not delayed.

The minimal initial flight segment of a telerobot could consist (in the subsystem breakdown described above) of all or parts of the MCM (Manipulator Control Mechanization) and S&P (Sensing and Perception) subsystems, since these are the subsystems that directly interact with the environment. The MCM subsystem is configured as two distinct elements, the real-time part and the non-real-time part. Clearly the real-time part, which performs position and force servoing of the manipulators, needs to fly with the telerobot. If very slow manipulation speeds are acceptable (and they may be dictated by safety concerns anyway), one or a few MIPS (Million Instructions Per Second) and 300 KFLOPS (Floating Point Operations per Second) will probably be adequate (a microVAX-II at about 0.75 MIPS and 100 KFLOPS has been used in our recent demonstrations of telerobotic capability, which included grappling and docking with a satellite mock-up, door opening, crank turning, etc.). If flexible arms are used or faster manipulations (10's of cm/sec) are needed (requiring dynamics computations) then perhaps 10-20 MIPS and 1-10 MFLOPS will be needed for arm control.

Of the Sensing and Perception subsystem, initially one could fly just the video cameras. Everything that is currently done in the tesbed S&P subsystem could be done on the ground if necessary, including the real-time edge detection used for grappling the satellite. This results from the fact that angular momentum is conserved with such high accuracy in orbit that the computation of the position, velocity, orientation, and angular-velocity on the ground can be very accurately projected a few seconds ahead in time. This requires accurate knowledge of the inertia tensor of the satellite, which can be computed from the CAD data base used in manufacture or derived from a few extra minutes of tracking the satellite. Once the satellite is grappled and docked, all current S&P computations are verifications of static object positions, and so the time delay will not seriously affect overall system performance. Likewise the Run Time Controller and Artificial Intelligence Planner can be located on the ground, as the small amount and low frequency of communication with the flight segment of the telerobot, and its non-time-critical nature is exemplified by our use of the ethernet for communication among these subsystems. Also, the non-real-time portion of the MCM subsystem (the trajectory generator) does not really need to fly, since the form of uplink data can be intermediate trajectory points or spline-fit parameters for free-space motion, task frame definitions for compliant motion (e.g. those axes in cartesian space that should be position controlled and those that should be force or torque controlled), the expected force/torque envelopes, and error conditions for abort.

If it becomes necessary to fly more of the S&P subsystem (for example, to assemble a large trusswork that is constantly vibrating), one possible approach being researched at JPL is for a multi-resolution pyramid image processor to be flight qualified. This processor, which filters and subsamples images to form a pyramid of low-pass and band-pass images (e.g. 512x512, 256x256, 128x128, etc.), takes about 200 MIPS to continuously compute the pyramid on a 30 Hz image stream at 512x512 resolution. This pyramid machine can be built using 3x3 convolver chips developed by JPL for a machine vision research tool, and which have been designed to be readily flight-qualified. This would permit real-time tracking, object acquisition, and reflex actions.

Other partitions of the flight and ground systems are possible. However, it does not seem advisable to partition S&P except at the camera outputs or after the feature extraction step due to the huge amount of parallel data being processed during feature extraction (some 10 parallel image paths at 10 Mpixels/sec for a total of about one Gbit/sec), which would swamp the available communications channels. Once one has gone to all the trouble of doing the feature extraction (~1 billion 12-bit fixed point operations per second), we might as well do all the rest of the S&P computations at the remote site (10-30 MIPS and 3-30 MFLOPS). RTC may take 10-100 MIPS and 10-50 MFLOPS as well, and is a good candidate for a flight hypercube or other concurrent architecture. As mentioned before, MCM could take up to 10-20 MIPS and 1-10 MFLOPS if manipulation speeds call for dynamics computations or if somewhat flexible arms are used.

The ground segment of the minimal flight telerobot need not be very complex to be highly effective. The recent demonstrations of telerobotic capability at JPL allowed the operator to use a wire-frame overlay on the video image returned from the remote site to designate objects. More recently, RCA, under contract to JPL to create an integrated Operator Control Station for the telerobot testbed, has demonstrated the use of wire-frames for 'analogic' designation in conjunction with voice input and output for 'symbolic' designation. One can certainly imagine an operator using the two six-degree-of-freedom hand controllers to control the position and orientation of wire-frame overlays while using voice control to call up named objects or give new names to object models being created. The operator, by matching the position of the wire-frames on the unprocessed video from the telerobot site in a stereo display, will be able to achieve millimeter precision in 'telling' the control system where objects are located (machine vision techniques can refine this further if needed). The operator can then invoke 'skills' such as bolt removal, handle grasping, or other operations involving contact with the environment using a combination of analogic 'pointing' and symbolic voice inputs. A repertoire of a few dozen 'skills' will allow an operator to conduct a complex task without any significant advance preparation or extensive data base, with or without a time delay in the control loop. If a complete task data base is available, then a 'virtual force field' can be created around objects in the environment and used to backdrive the FRHC's to avoid contact while teleoperation is underway. The only physical contact permitted by the system would be when control is 'traded' to the autonomous system for execution of a skill, or when 'shared' control is invoked, where the autonomous system controls all forces while the human directs motion in unconstrained directions. Demonstration of this shared and traded control methodology is a principal objective for the telerobot testbed in the coming year.

CONCLUSIONS

A near-term telerobot flight segment needs only 1) video cameras and 2) the inner core of the servo-control system, able to maintain stable control over the arms while in free space motion, to decellerate smoothly near the task, to move very slowly in a guarded move to the instant of contact, and to switch to force control for executing the appropriate compliant frame definition for the task at hand. A processor with as little as 2 MIPS and 300 KFLOPS should be adequate for this minimal time-delayed telerobot. The downlink would be video at a few hundred Mbits per second, and the uplink would be trajectory spline parameters, compliant frame definitions, and error condition predicates at ~100 Kbits per second.

In the longer term it may become advantageous to fly the entire MCM, S&P, RTC, and part of the AIP subsystems (some of the Artificial Intelligence Planner will always be located at the operator conrol station for replanning and to give advice). This, however, will only improve the performance by an order-of-magnitude or so over the minimal system described above (which has some four to five orders-of-magnitude economic advantage over the use of astronauts for these tasks). The computational load of the flight subsystems will jump from a few MIPS and some 0.3 MFLOPS to perhaps 30-100 MIPS and 10-50 MFLOPS of general-purpose computer and 1000's of MIPS of image processor. Thus great advances in flight-qualified computation are needed to evolve much beyond the minimal time-delayed telerobot. Yet the minimal time-delayed telerobot, which can be configured from the imminently flight-qualified NS32016 or 80386 processor families, offers huge economic benefits for orbital assembly, maintenance, and repair tasks.

ACKNOWLEDGMENT

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

REFERENCES

1. T.B. Sheridan and W.R. Ferrell, 'Remote Manipulative Control with Transmission Delay,' *IEEE Transactions on Human Factors in Electronics*, p 25-29 (September 1963).
2. T.B. Sheridan and W.R. Ferrell, *Man-Machine Systems: Information, Control, and Decision Models of Human Performance*, MIT Press, Cambridge MA, 1974.
3. W. F. Zimmerman and J. R. Matijevic, 'System Engineering Techniques for Establishing Balanced Design and Performance Gidelines for the Advanced Telerobotic Testbed,' *Proceedings of the Workshop on Space Telerobotics*, JPL publication 87-13, vol 1,p 67-74.
4. B. Wilcox, D.B. Gennery, B. Bon, and T. Litwin, 'The Sensing and Perception Subsytem of the NASA Research Telerobot,' *ibid*, Vol 2, p 3-8.
5. S. Hayati and B. Wilcox, 'Manipulator Control and Mechanization: A Telerobot Subsytem', *ibid*, Vol 2, p 219-234.
6. J. Balaram, A. Lokshin, K. Kreutz, and J. Beahan, 'A Run-Time Control Architecture for the JPL Telerobot', *ibid*, Vol 1, p 211-222.
7. C. Collins and M. Rokey, 'Planning for the JPL telerobotics project,' Proc. *4th Annual Artificial Intelligence and Advancd Technology Conference*, p 429-437, May 1988.
8. D. Mittman, 'Audrey: An Interactive Simulation and Spatial Planning Environment for the NASA Telerobot System,' *ibid*, p 421-428.
9. A. K. Bejczy and Z. Szakaly, 'Universal Computer Control System (UCCS) for Space Telerobots', *IEEE International Conference on Robotics and Automation,*, March-April 1987, vol 1, p318-325.